Sequence-to-point learning with neural networks for non-intrusive load monitoring

Chaoyun Zhang, Mingjun Zhong, Zongzuo Wang, Nigel Goddard, & Charles Sutton
School of Informatics, University of Edinburgh, United Kingdom
chaoyun.zhang@ed.ac.uk, {mzhong,ngoddard,csutton}@inf.ed.ac.uk

arXiv:1612.09106v2 [stat.AP] 24 Feb 2017

Abstract

Energy disaggregation (a.k.a nonintrusive load monitoring, NILM), a single-channel blind source separation problem, aims to decompose the mains which records the whole electricity consumption into appliance-wise readings. This problem is difficult because it is inherently unidentifiable. Recent approaches have shown that the identifiability problem could be reduced by introducing domain knowledge into the model. Deep neural networks have been shown to be promising to tackle this problem in literatures. However, it is not clear why and how the neural networks could work for this problem. In this paper, we propose sequence-to-point learning for NILM, where the input is a window of the mains and the output is a single point of the target appliance. We use convolutional neural networks to train the model. Interestingly, we systematically show that the convolutional neural networks can inherently learn the signatures of the target appliances, which are automatically added into the model to reduce the identifiability problem. We applied the proposed neural network approaches to a real-world household energy data, and show that the methods achieve state-of-the-art performance.

1 Introduction

Energy disaggregation (a.k.a nonintrusive load monitoring, NILM) [Hart, 1992] is a single-channel blind source separation (BSS) problem that aims to decompose the whole energy consumption of a building into the energy usage of individual appliances. The purpose is to help households to reduce their energy consumption by helping them to understand what is causing them to use energy, and it has been shown that disaggregated information can help householders to reduce energy consumption by as much as $5-15\%$ [Fischer, 2008]. However, current electricity meters can only report the whole-home consumption data. This triggers the demand of machine-learning tools to infer the appliance-specific consumption.

Energy disaggregation is unidentifiable and thus a difficult prediction problem because it is a single-channel BSS problem; we want to extract more than one source from a single observation. Additionally, there are a large number of sources of uncertainty in the prediction problem, including noise in the data, lack of knowledge of the true power usage for every applicant in a given household, multiple devices exhibiting similar power consumption, and simultaneous switching on/off of multiple devices. Therefore energy disaggregation has been an active area for the application of artificial intelligence and machine learning techniques. Popular approaches have been based on factorial hidden Markov models (FHMM) [Kolter and Jaakkola, 2012; Parson *et al.*, 2012; Zhong *et al.*, 2013; Zhong *et al.*, 2014; Zhong *et al.*, 2015; Lange and Bergés, 2016] and signal processing methods [Pattem, 2012; Zhao *et al.*, 2015; Zhao *et al.*, 2016; Batra *et al.*, 2016; Tabatabaei *et al.*, 2017]. To tackle the identifiability problem, various approaches have been proposed by incorporating domain knowledge into the model. For example, local information, e.g., appliance power levels, ON-OFF state changes, and durations, was incorporated into the model [Kolter and Jaakkola, 2012; Parson *et al.*, 2012; Pattem, 2012; Zhao *et al.*, 2015; Tabatabaei *et al.*, 2017]; others are interested in incorporating global information, e.g., total number of cycles and total energy consumption [Zhong *et al.*, 2014; Zhong *et al.*, 2015; Batra *et al.*, 2016]. However, the domain knowledge required by these methods needs to be extracted manually, which makes the methods more difficult to use.

Recently, it has been shown that single-channel BSS could be modelled by using sequence-to-sequence learning with neural networks [Grais *et al.*, 2014; Huang *et al.*, 2014; Du *et al.*, 2016]. Particularly, it has been applied to energy disaggregation [Kelly and Knottenbelt, 2015a] —both convolutional (CNN) and recurrent neural networks (RNN) were employed. The idea of sequence-to-sequence learning is to treat a window of the mains as the input of a neural network and the corresponding window of appliance power readings as the target. For simplicity, we name this approach as SeqNet. From the view of probabilistic modelling, the power level of each appliance is assumed to be drawn from a Gaussian distribution where the mean is the output of the neural network and the variance is a constant. Despite promising results, this approach has several difficulties: 1) Each appliance power level at a single time step is predicted many times, once for each sliding window; an average of multiple predictions is naturally used, which consequently smooths the edges; 2) A regression model is used for modelling the tar-

gets, which results in a deterministic prediction; this does not take into account the uncertainty in the predictions; 3) Neural networks are black boxes and it is difficult to understand exactly what information they have learnt. To have confidence in the models, it is vital to interpret the model predictions and understand what information the neural networks for NILM are relying on to make their predictions.

In this paper, we propose a different architecture called sequence-to-point learning for single-channel BSS. Given a window of the mains sequence as the input, instead of treating the corresponding window of the target appliance as the outputs, we only want to use the single midpoint of the appliance window as the output. We propose two different models: 1) the midpoint is represented as a non-linear regression of the mains window (which is denoted by PointNet); 2) the midpoint is quantized and therefore the problem could be transformed into a classification problem (which is denoted by ProbNet). In our experiments, we employed CNN for our models. We also implemented RNN, but we found that CNN performed the best. Besides, it is well known that CNN has lower computational costs than RNN.

Importantly, we experimentally analysed what the neural networks have learnt by visualizing the feature maps of CNN. We found that the neural networks have automatically extracted useful signatures from the mains window. Particularly, those meaningful signatures are the following: 1) the change points of the time series which indicate turning ON or OFF of the appliance; 2) the amplitude of the signal which indicates the power level of the appliance; 3) the length of the steady state of the time series which indicates the duration of a state of the appliance. Interestingly, we have noted that in the literature these signatures were the key constraints added to the models, but it had to resort to handcrafted feature selection procedures [Kolter and Jaakkola, 2012; Parson *et al.*, 2012; Pattem, 2012; Zhao *et al.*, 2015; Zhong *et al.*, 2014; Zhong *et al.*, 2015; Batra *et al.*, 2016; Tabatabaei *et al.*, 2017].

We compared these methods by applying them to the UK-DALE data set, which measures the electricity usage of five UK households [Kelly and Knottenbelt, 2015b]. We find that sequence-to-point learning performs dramatically better than previous work, with as much as 68% reduction in error compared to previously published deep learning approaches based on sequence-to-sequence learning.

2 Energy disaggregation

The goal of energy disaggregation is to recover the energy consumption of individual appliances from the mains. Suppose we have observed the mains Y which indicates the total power in Watts in a household, where $Y = (y_1, y_2, ..., y_T)$ and $y_t \in R_+$. Suppose there are a number of appliances in the same house. For each appliance, its reading is denoted by $X_i = (x_{i1}, x_{i2}, ..., x_{iT})$, where $x_{it} \in R_+$. At each time step, y_t is assumed to be the sum of the readings of individual appliances, possibly plus a Gaussian noise factor with zero mean and variance σ^2 such that $y_t = \sum_i x_{it} + \epsilon_t$. Often we are only interested in I appliances, i.e., the ones that use the most energy; others will be regarded as an unknown factor $\mathbf{u} = (u_1, \cdots, u_T)$. The model could then be represented as

$$y_t = \sum_{i=1}^{I} x_{it} + u_t + \epsilon_t. \quad (1)$$

The prior of $\mathbf{u}$ is assumed to be $p(\mathbf{u}) \propto \prod_{t=1}^{T-1} \exp(\lambda|u_{t+1} - u_t|)$. The task of energy disaggregation is to infer X_i given Y.

2.1 Additive factorial hidden Markov model

The additive factorial hidden Markov model (AFHMM) is a natural approach to represent the model (1) [Kolter and Jaakkola, 2012; Pattem, 2012; Zhong *et al.*, 2014]. AFHMM assumes that $\mathbf{x}_i$ is a hidden Markov chain with states $Z_{it} \in \{1, 2, ..., K_i\}$. In AFHMM, the model parameters θ are assumed to be estimated by using observation data. These parameters are the initial probabilities $\pi_i = (\pi_{i1}, ..., \pi_{iK_i})$, the mean values $\mu_i = (\mu_{i1}, \cdots, \mu_{iK_i})$ such that $x_{it} \in \mu_i$, and the transition probabilities $p^i_{j,k} = P(Z_{it} = j|Z_{i,t-1} = k)$. To infer x_i is now equivalent to infer the most probable states Z_{it}. Instead of using Z, we could now use a binary vector $S_{it} = (S_{it1}, \cdots, S_{itK_i})^T$ to represent Z such that $S_{itj} = 1$ when $Z_{it} = j$ and for all $S_{itk} = 0$ when $k \neq j$. The posterior distribution of AFHMM has the form

$$P(S|Y) \propto \prod_{t=1}^{T} p(Y_t|S_t) \prod_{i=1}^{I} P(S_{i1}) \prod_{t=2}^{T} \prod_{i=1}^{I} P(S_{it}|S_{i,t-1}),$$

where we assume $p(Y_t|S_t) \sim \mathcal{N}(\sum_{i=1}^{I} \mu_i^T S_{it} + u_t, \sigma_t^2)$. The goal is to infer the states S.

Various inference algorithms could then be employed to infer S [Kolter and Jaakkola, 2012; Zhong *et al.*, 2014; Shaloudegi *et al.*, 2016]. It is well-known that the problem is still unidentifiable. Therefore, different approaches have been proposed to use domain knowledge to reduce this problem [Zhong *et al.*, 2014; Zhong *et al.*, 2015]. However, as it was noted, all these approaches require handcrafted features based on the observation data. Instead, we propose to employ neural networks to extract those features automatically during learning.

2.2 Sequence-to-sequence learning

Kelly and Knottenbelt [2015a] proposed to apply deep learning methods to NILM. The model learns a nonlinear regression between the sequences of the mains and appliances. A sequence of the mains is the input and the sequence of target appliance with the same time stamps is the target of the neural networks, which is a sequence-to-sequence learning method. Both CNNs and RNNs were used. Given a period of aggregated data, their models are able to predict the appliance-wise energy usage in the some period. For notation simplicity, we call this model SeqNet. It has been shown that SeqNet outperforms the baseline AFHMM. However it is not yet known why and how the neural networks work for energy disaggregation. It should be noted that similar deep neural network approaches have been applied to the single-channel BSS problems in audio and speech [Grais *et al.*, 2014; Huang *et al.*, 2014; Du *et al.*, 2016].

3 Sequence-to-point learning for energy disaggregation

Instead of training a network to predict a window of appliance readings, we propose to train a neural network to only predict the midpoint of that window. The idea is that the input of the network is the mains window, and the output is the midpoint of the corresponding window of the target appliance. We call this type of method a sequence-to-point learning method, which is widely applied for modelling the distributions of speech and image [van den Oord *et al.*, 2016a; van den Oord *et al.*, 2016b]. There are two different ways to model the midpoint. Firstly, the midpoint is represented as a non-linear regression of the mains window (denoted by PointNet); this method assumes that the appliance readings are continuous. However, in the real world, all the electricity readings tend to be stored in integer with a lower unit, e.g., Joule. It has at least two advantages —using a lower unit makes the reading to be more accurate and integers need less disk space. Based on this fact, every appliance has a finite number of states. Therefore, alternatively the midpoints could be represented as classification outputs of a neural network with the mains window being the input (denoted by ProbNet).

In general, given a mains window $Y = (y_1, y_2, \cdots, y_T)$, we are interested in inferring the possible readings of an appliance which are $X = (x_1, x_2, \cdots, x_T)$. We are thus interested in the posterior distribution $p(X|Y)$, which is represented as a single neural network in SeqNet. For the sequence-to-point learning, we assume the posterior distribution is conditionally independent such that

$$p(X|Y) = \prod_{t=1}^{T} p(x_t|Y). \tag{2}$$

From the view of probabilistic modelling, PointNet assumes that $p(x_t|Y)$ is a Normal distribution $x_t \sim \mathcal{N}(\mu_\theta(Y), \sigma^2_\theta(Y))$. Both the mean and variance could be represented as a neural network, and so the model parameters θ could be trained by using the training data. In this paper, we assume $\sigma^2_\theta(Y) = 1$. ProbNet assumes that $p(x_t|Y)$ is a softmax distribution such that $P(x_t = i) = \sigma(\mu_i)$ where $\sigma(\mu_i) = \frac{\exp(\mu_i)}{\sum_{j=1}^{K}\exp(\mu_j)}$. The softmax distribution which assumes a discrete distribution has been applied in image problems [van den Oord *et al.*, 2016b].

In order to use a softmax distribution, we quantize the appliance readings since there are a large number of possible values for every appliance. For example, the maximum power of a kettle is over 3,000 Watts, which leads to a dense connection between softmax layer and the last hidden layer. In order to reduce the number of states, we firstly map x to the range $[0, 1]$ by dividing by the maximum power (see Table 1); secondly, we transform these values via μ-law companding transformation [Brokish and Lewis, 1997]

$$y = \text{sign}(x)\frac{\ln(1+\mu|x|)}{\ln(1+\mu)} \tag{3}$$

where $\mu = 256$; finally, we quantize them into 256 possible states. To recover the original values, we apply the reverse order of the procedure. The inverse equation of μ-law transformation is $x = \text{sign}(y)(1/\mu)((1+\mu)^{|y|} - 1)$.

3.1 Architectures

For the sequence-to-sequence learning, we employed the denoising autoencoders described in Kelly and Knottenbelt [2015a], because this performed better than the other architectures that they proposed. For the sequence-to-point learning, we use the convolutional neural network. The architectures are shown in Figure 1. Note that we also tried the mixture of CNN and RNN, however, the performance were not improved over CNN, so we do not report them here. One reason behind this observation might be that the local features learnt by the CNN probably dominates the disaggregation performance; this point will be experimentally analysed in Section 5.5.

4 Integrating neural networks and AFHMM

The distributions estimated by neural networks and AFHMM can be integrated by using the logarithmic opinion pool method Bordley [1982]. It essentially seeks a new distribution $\widetilde{P}(X|Y)$ that weights the opinions from $P_A(X|Y)$ and $P_N(X|Y)$, where P_A and P_N denotes the distributions estimated by AFHMM and neural networks respectively. $\widetilde{P}$ can be obtained by using the KL-divergence,

$$\begin{aligned}\widetilde{P}(X|Y) \equiv \text{argmin}_{\widetilde{P}} \quad & [\alpha_1 KL(\widetilde{P}(X|Y)|P_A(X|Y)) \\ & +\alpha_2 KL(\widetilde{P}(X|Y)|P_N(X|Y))],\end{aligned}$$

where $\alpha_1, \alpha_2 > 0$ and $\alpha_1 + \alpha_2 = 1$. The α_1 and α_2 are viewed as the weights of distributions. The optimal solution has the form $\widetilde{P}(X|Y) \propto P_A(X|Y)^{\alpha_1} P_N(X|Y)^{\alpha_2}$. An optimal $\widetilde{X}$ could be inferred by maximizing the posterior $\widetilde{P}$. We set the weights α_1, α_2 to be 0.5 in this paper. We convert this problem into a convex quadratic program by relaxing the variables as in [Zhong *et al.*, 2014], which is efficiently solved by using the solver MOSEK [ApS, 2015].

5 Experiments

We implemented five models on the same dataset for energy disaggregation, which are AFHMM [Kolter and Jaakkola, 2012], latent Bayesian melding (LBM) [Zhong *et al.*, 2015], SeqNet [Kelly and Knottenbelt, 2015a], PointNet, and ProbNet. To verify the effectiveness and efficiency, we conduct comprehensive comparisons in terms of different performance metrics and model complexity. The deep learning models are implemented in Python using TensorFlow. The networks were trained on machines with NVIDIA GTX 970 and NVIDIA GTX TITAN X GPUs.

5.1 UK-DALE dataset

We report results on the UK-DALE dataset [Kelly and Knottenbelt, 2015b], which measured the domestic appliance-level energy consumption and whole-house energy usage of five UK houses. All the readings were recorded in every 6 seconds from November 2012 to January 2015. The dataset contains the measurements of over 10 types of appliances,

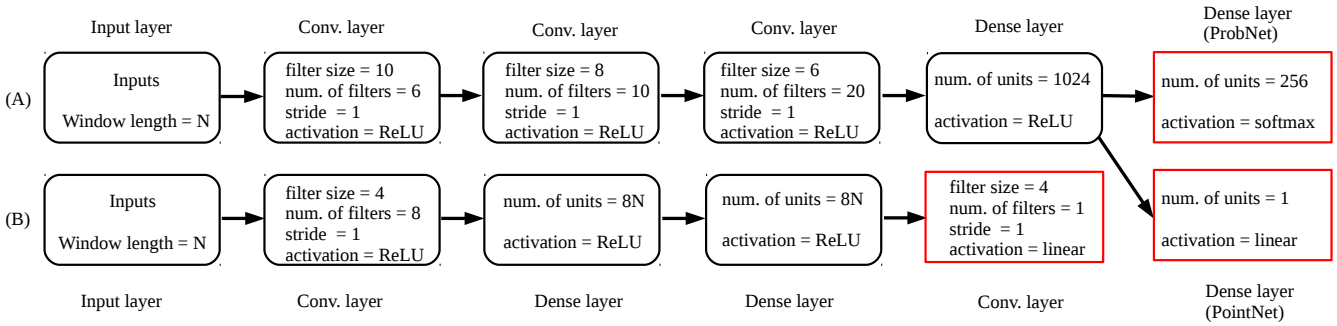

Figure 1: The architectures for (A) the sequence-to-point and (B) the sequence-to-sequence neural networks.

Table 1: The parameters used in this paper for each appliance. Power unit is Watt.

	Kettle	Microwave	Fridge	Dish Washer	Washing Machine
Window length (point)	129	129	299	599	599
Maximum power	3948	3138	2572	3230	3962
On power threshold	2000	200	50	10	20
Mean on power	700	500	200	700	400
Standard deviation on power	1000	800	400	1000	700

however, in this paper we are only interested in kettle, microwave, fridge, dish washer and washing machine which are the popular appliances to be used for assessing NILM algorithms. We used the houses 1, 3, 4, and 5 for training the neural networks, and house 2 as the test data. Note that we are therefore considering the transfer learning setting in which we train and test on different households. This setting has the challenge that the same type of appliance will vary in its power demands in different houses, but good performance in the transfer learning set-up is vital to practical application of NILM methods.

5.2 Data preprocessing

We describe how the training data were prepared for training the sequence-to-point neural network. A window of the mains was used as the input sequence; the window length for each appliance is shown in the Table 1. The training windows were obtained by sliding the mains reading at every point; the midpoints of the corresponding appliance were used as the targets. For PointNet, both the input windows and targets were preprocessed by subtracting the mean values and dividing by the standard deviations (see these parameters in the Table 1); for ProbNet, the difference is that the targets were preprocessed by the quantization procedure described in Section 3 using the maximum power parameters shown in the Table 1.

5.3 Performance evaluation

We use three metrics to compare these approaches. Denote x_t and $\hat{x}_t$ as the ground truth and prediction of an appliance at time t, respectively. The NILM performance can be evaluated by the normalized disaggregation error (NDE) [Zhong *et al.*, 2014]

$$\text{NDE} = \frac{\sum_t (\hat{x}_t - x_t)^2}{\sum_t x_t^2}, \tag{4}$$

which is invariant to the unit. When we are interested in the error in power at every time point, we use the mean absolute error (MAE)

$$\text{MAE} = \frac{1}{T}\sum_{t=1}^{T} |\hat{x}_t - x_t|. \tag{5}$$

This provides a measure of errors that is less affected by outliers, i.e. isolated predictions that are particularly inaccurate. When we are interested in the total error in energy over a period, in this case, one day, we use the normalised signal aggregate error (SAE)

$$\text{SAE} = \frac{|\hat{r} - r|}{r}, \tag{6}$$

where r and $\hat{r}$ denote the ground truth and inferred total energy consumption of an appliance, that is $r = \sum_t x_t$ and $\hat{r} = \sum_t \hat{x}_t$. This measure is useful because a method could be accurate enough for reports of daily power usage even if its per-timestep prediction is less accurate.

5.4 Experimental results

For comparison purposes, we applied these methods to test on one-month data from house 2; then computed these evaluation measures. All of the results are shown in Table 2, 3, and 4. We also applied the latent Bayesian melding (LBM) approach to the same data which uses global information of the appliances [Zhong *et al.*, 2015].

All the other methods outperform AFHMM which is considered as the baseline. The PointNet performed the best comparing to all the other methods. ProbNet outperformed SeqNet in terms of MAE and SAE, but SeqNet performed better in terms of NDE. We note that SeqNet has higher computational costs, since the number of parameters to train in

Table 2: The appliance-level normalized disaggregation error (NDE) (Watt). Best results are shown in bold.

Methods	Kettle	Microwave	Fridge	Dish w.	Washing m.	Overall
AFHMM [Kolter and Jaakkola, 2012]	1.32	2.80	1.00	1.71	8.74	3.11 ± 2.87
LBM [Zhong *et al.*, 2015]	0.41	1.14	0.42	0.43	2.17	0.91 ± 0.68
SeqNet [Kelly and Knottenbelt, 2015a]	**0.10**	**0.34**	0.46	0.48	1.19	0.51 ± 0.36
PointNet (This paper)	**0.10**	0.69	**0.27**	**0.19**	**0.27**	$\mathbf{0.30 \pm 0.20}$
ProbNet (This paper)	0.35	0.97	0.42	0.62	0.57	0.58 ± 0.21

Table 3: The appliance-level mean absolute error (MAE) (Watt). Best results are shown in bold.

Methods	Kettle	Microwave	Fridge	Dish w.	Washing m.	Overall
AFHMM [Kolter and Jaakkola, 2012]	47.38	21.18	42.35	199.84	103.24	82.79 ± 64.50
LBM [Zhong *et al.*, 2015]	19.51	9.62	26.45	54.62	47.23	31.48 ± 16.91
SeqNet [Kelly and Knottenbelt, 2015a]	9.08	9.33	22.80	28.56	27.63	19.47 ± 8.61
PointNet (This paper)	**7.96**	8.42	**17.20**	**18.03**	**8.94**	$\mathbf{12.11 \pm 4.51}$
ProbNet (This paper)	10.68	**7.95**	19.04	28.85	**8.88**	15.08 ± 7.92

Table 4: The appliance-level signal aggregate error (SAE) (Watt). Best results are shown in bold.

Methods	Kettle	Microwave	Fridge	Dish w.	Washing m.	Overall
AFHMM [Kolter and Jaakkola, 2012]	1.06	1.04	0.98	4.50	8.28	3.17 ± 2.88
LBM [Zhong *et al.*, 2015]	0.40	0.51	**0.07**	0.53	2.94	0.89 ± 1.03
SeqNet [Kelly and Knottenbelt, 2015a]	0.11	0.79	0.34	0.33	1.33	0.58 ± 0.43
PointNet (This paper)	**0.07**	**0.23**	0.18	**0.26**	**0.16**	$\mathbf{0.18 \pm 0.06}$
ProbNet (This paper)	0.37	0.75	0.28	0.53	0.59	0.50 ± 0.16

SeqNet is 20 times more than ProbNet and PointNet (see Table 5). Both AFHMM and LBM, which use handcrafted features, were no better than the neural network approaches. This indicates that neural networks can extract more meaningful features from the data for representing the target appliances (see Section 5.5). Overall, we indicate that the sequence-to-point model has not only improved the performance, but also reduced the computational costs over the sequence-to-sequence model. We show some disaggregation results in Figures 3-7.

Table 5: Number of parameters of the networks (averaged over the individual networks for each appliance).

	Average number of parameters
SeqNet	172,925,496
PointNet	7,192,355
ProbNet	7,453,680

5.5 Visualization of latent features

We have shown that the neural networks approaches outperform the hidden Markov models. The question is what the neural networks have learnt from the data? Presumably, we expect that a good model should be able to recover the appliances if we could tell the model the signatures of the appliances. For example, we could tell the model these information: 1) a kettle only has two states: ON and OFF, and when it is ON the power should be approximately $2,000 - 3,000$ Watts; 2) the approximate duration of the kettle when it is ON. This information could be enough to detect a kettle. Actually, these information (or signatures) could be described by summary statistics. It has also been shown that it greatly improved the performance of the algorithms [Zhao *et al.*, 2015; Batra *et al.*, 2016; Zhong *et al.*, 2015]. Could we define a model which automatically learns these signatures and also at the same time is able to recover the appliances?

Interestingly, we observed that sequence-to-point leaning is inherently learning these signatures. We used ProbNet to see those meaningful signatures; we systemically change the input of the network to see what has been learnt. In these experiments we looked at the kettle which is easier to study because there are less number of states. For each different input, we plotted the feature maps of the last CNN layer in the Figure 2. It is interesting that all the filters detected the state changes of the appliances. More specifically, some filters take the responsibility of detecting amplitude of the appliance and as well as the state changes, but others only detect the state changes. Figure 2 (b) shows that when the kettle was manually removed, the network thinks that the amplitude of the signal and as well as the duration were not appropriate for a kettle. Figure 2 (c) shows that when the amplitude of the kettle was double, the network detects the kettle which is reasonable because both the duration and amplitude correspond to a kettle. Figure 2 (d) indicates that when the amplitude of the kettle was manually reduced, the network thinks there was no kettle. Figure 2 (e) shows that when the duration of the appliance usage was set too long (> 8 minutes), the network thinks it was not a kettle. Figure 2 (f) shows that the network works when there was no kettle.

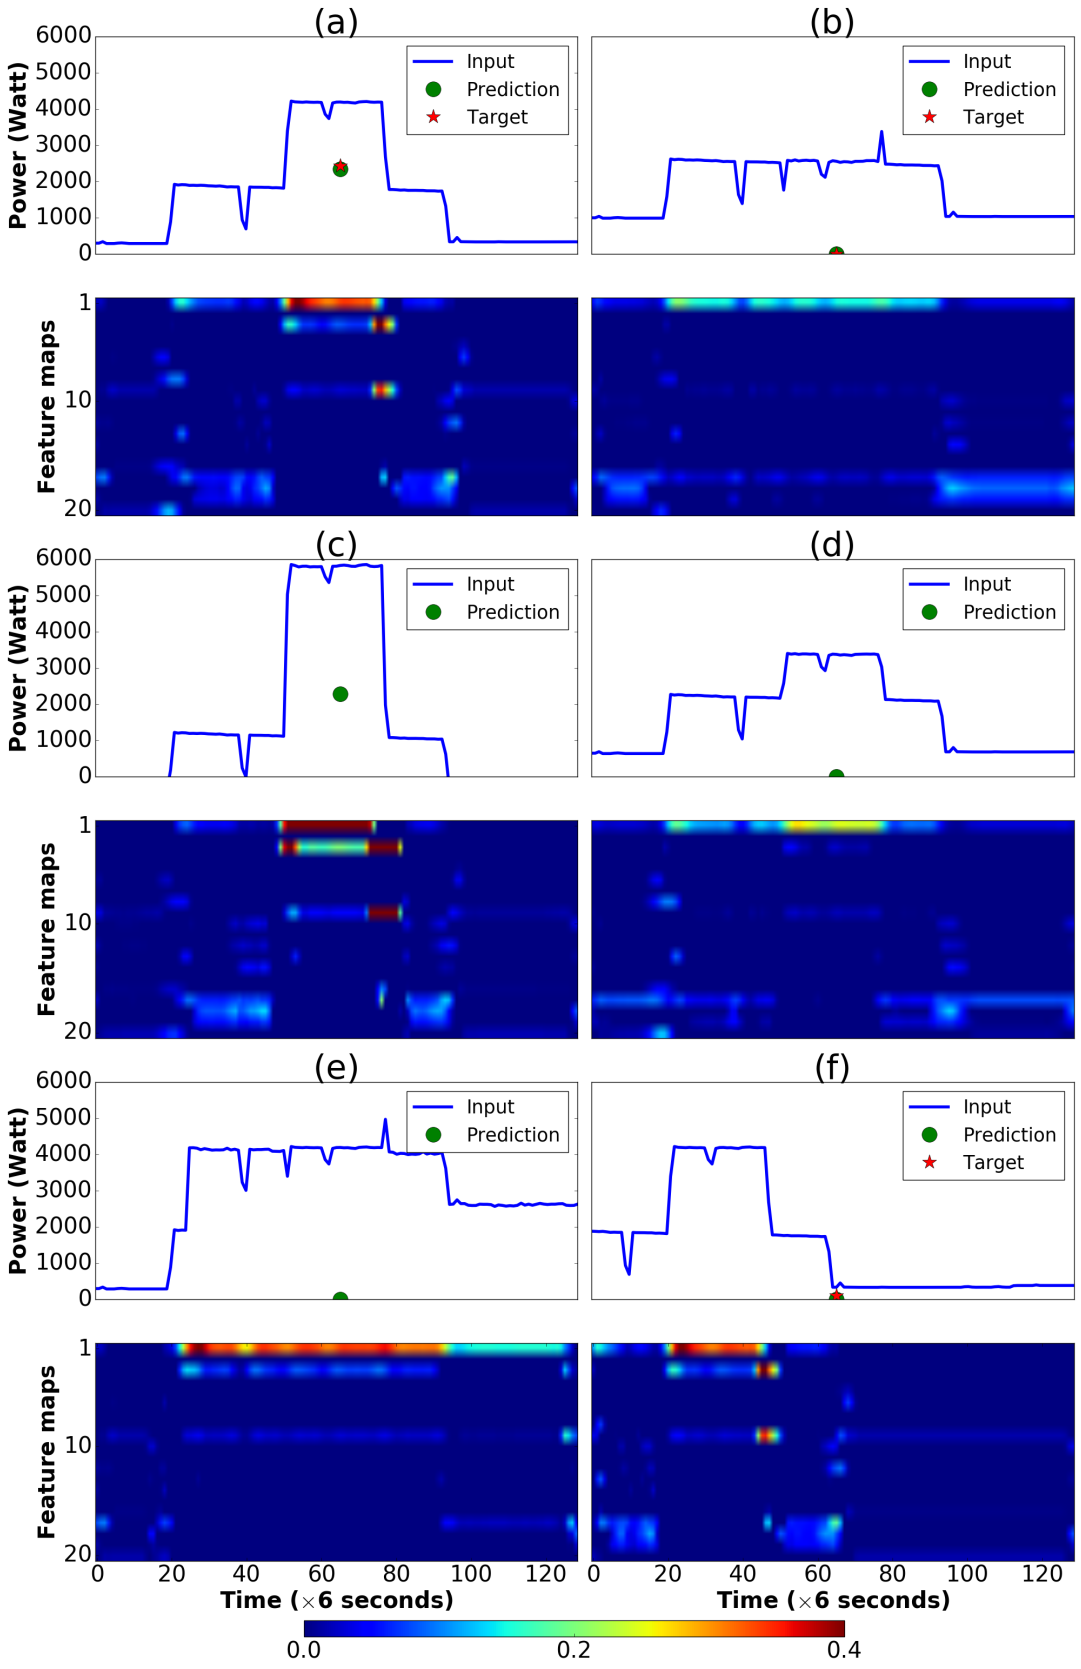

Figure 2: Feature maps learnt by the convolutional networks for various types of inputs (the mains). The feature maps contain the signatures of an appliance, which are used to predict the states of the appliance. These plots indicate that the learnt signatures are when an appliance is turned on and off (see the edges of the feature maps), the duration of an appliance when it is turned on, and the power level (red indicates higher power level). (a) The kettle is in the mains; the network correctly predicted the target. (b) The kettle was manually removed from the mains. (c) The power level of the kettle was set be double of the true level. (d) The power level of the kettle was set to be half of the true level. (e) The duration of the kettle was set to be double. (f) Target is at the edge of the mains reading.

6 Conclusions

In this paper, we have proposed a sequence-to-point learning with neural networks for energy disaggregation. We have applied the proposed schemes to a real world data set. We have shown that sequence-to-point learning outperforms previous work using sequence-to-sequence learning in terms of performance and computational costs. By visualizing the learnt feature maps from neural networks, we have shown that the neural networks can learn inherently meaningful features from the data, which are crucial signatures for performing energy disaggregation. It would be interesting to apply the proposed methods to the single-channel blind source separation problems in other domains, for example, audio and speech.

References

[ApS, 2015] MOSEK ApS. *The MOSEK optimization toolbox for Python manual. Version 7.1 (Revision 60)*, 2015.

[Batra *et al.*, 2016] Nipun Batra, Amarjeet Singh, and Kamin Whitehouse. Gemello: Creating a detailed energy breakdown from just the monthly electricity bill. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 431–440. ACM, 2016.

[Bordley, 1982] Robert F. Bordley. A multiplicative formula for aggregating probability assessments. *Management Science*, 28(10):1137–1148, 1982.

[Brokish and Lewis, 1997] Charles W Brokish and Michele Lewis. A-law and mu-law companding implementations using the tms320c54x. *SPRA163*, 1997.

[Du *et al.*, 2016] Jun Du, Yanhui Tu, Li-Rong Dai, and Chin-Hui Lee. A regression approach to single-channel speech separation via high-resolution deep neural networks. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 24(8):1424–1437, August 2016.

[Fischer, 2008] Corinna Fischer. Feedback on household electricity consumption: a tool for saving energy? *Energy efficiency*, 1(1):79–104, 2008.

[Grais *et al.*, 2014] Emad M Grais, Mehmet Umut Sen, and Hakan Erdogan. Deep neural networks for single channel source separation. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3734–3738. IEEE, 2014.

[Hart, 1992] G.W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870 –1891, Dec 1992.

[Huang *et al.*, 2014] Po-Sen Huang, Minje Kim, Mark Hasegawa-Johnson, and Paris Smaragdis. Deep learning for monaural speech separation. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1562–1566. IEEE, 2014.

[Kelly and Knottenbelt, 2015a] Jack Kelly and William Knottenbelt. Neural NILM: Deep neural networks applied to energy disaggregation. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, pages 55–64. ACM, 2015.

[Kelly and Knottenbelt, 2015b] Jack Kelly and William Knottenbelt. The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. *Scientific Data*, 2(150007), 2015.

[Kolter and Jaakkola, 2012] Zico Kolter and Tommi S Jaakkola. Approximate inference in additive factorial hmms with application to energy disaggregation. In *AISTATS*, volume 22, pages 1472–1482, 2012.

[Lange and Bergés, 2016] Henning Lange and Mario Bergés. Efficient inference in dual-emission FHMM for energy disaggregation. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

[Parson *et al.*, 2012] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, pages 356–362, July 2012.

[Pattem, 2012] Sundeep Pattem. Unsupervised disaggregation for non-intrusive load monitoring. In *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*, volume 2, pages 515–520. IEEE, 2012.

[Shaloudegi *et al.*, 2016] Kiarash Shaloudegi, András György, Csaba Szepesvari, and Wilsun Xu. SDP relaxation with randomized rounding for energy disaggregation. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4979–4987. Curran Associates, Inc., 2016.

[Tabatabaei *et al.*, 2017] S. M. Tabatabaei, S. Dick, and W. Xu. Toward non-intrusive load monitoring via multi-label classification. *IEEE Transactions on Smart Grid*, 8(1):26–40, Jan 2017.

[van den Oord *et al.*, 2016a] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.

[van den Oord *et al.*, 2016b] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1747–1756, 2016.

[Zhao *et al.*, 2015] B. Zhao, L. Stankovic, and V. Stankovic. Blind non-intrusive appliance load monitoring using graph-based signal processing. In *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 68–72, Dec 2015.

[Zhao *et al.*, 2016] B. Zhao, L. Stankovic, and V. Stankovic. On a training-less solution for non-intrusive appliance load monitoring using graph signal processing. *IEEE Access*, 4:1784–1799, 2016.

[Zhong *et al.*, 2013] M. Zhong, N. Goddard, and C. Sutton. Interleaved factorial non-homogeneous hidden Markov models for energy disaggregation. In *Neural Information Processing Systems, Workshop on Machine Learning for Sustainability*, Lake Tahoe, Nevada, USA, 2013.

[Zhong *et al.*, 2014] Mingjun Zhong, Nigel Goddard, and Charles Sutton. Signal aggregate constraints in additive factorial HMMs, with application to energy disaggregation. In *Advances in Neural Information Processing Systems*, pages 3590–3598, 2014.

[Zhong *et al.*, 2015] Mingjun Zhong, Nigel Goddard, and Charles Sutton. Latent Bayesian melding for integrating individual and population models. In *Advances in Neural Information Processing Systems*, pages 3618–3626, 2015.

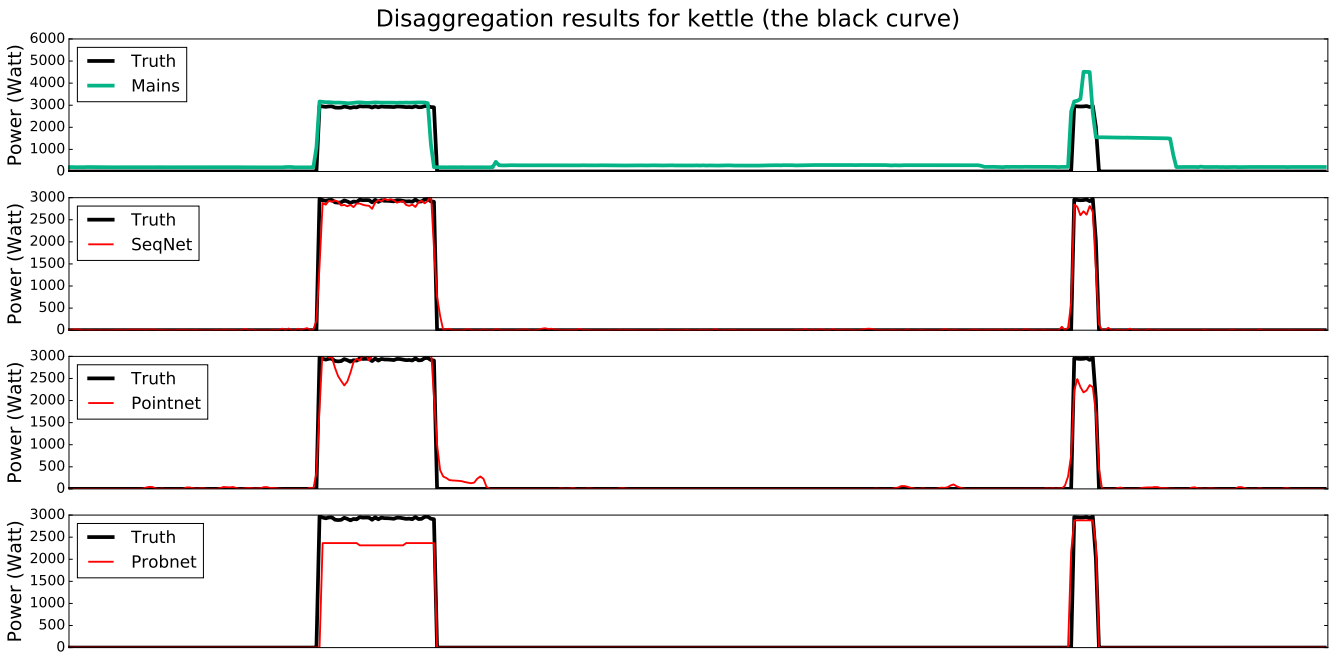

Figure 3: The example disaggregation results for kettle.

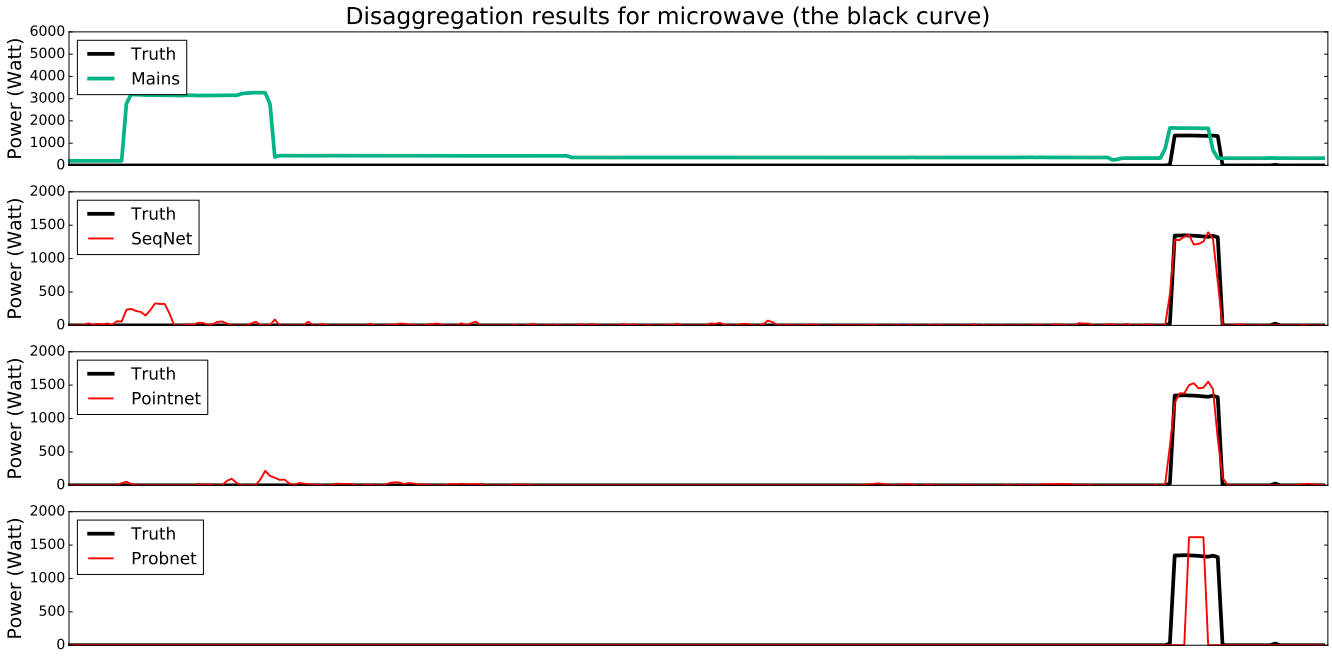

Figure 4: The example disaggregation results for microwave.

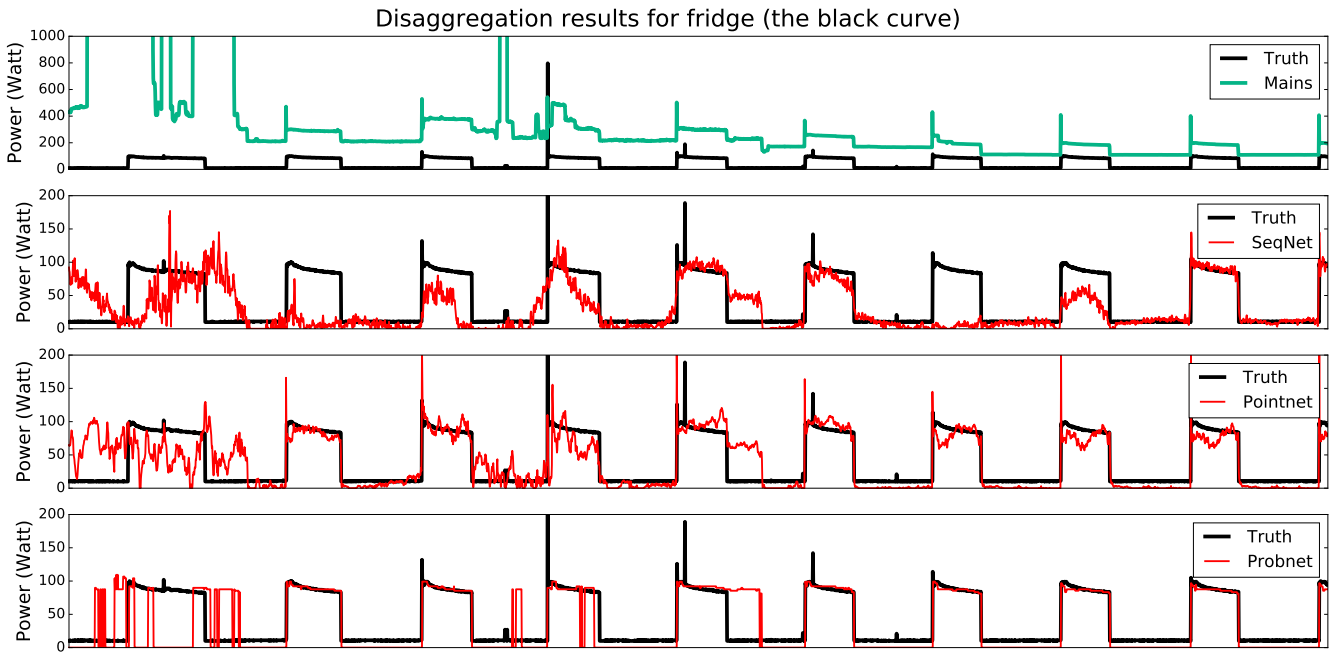

Figure 5: The example disaggregation results for fridge.

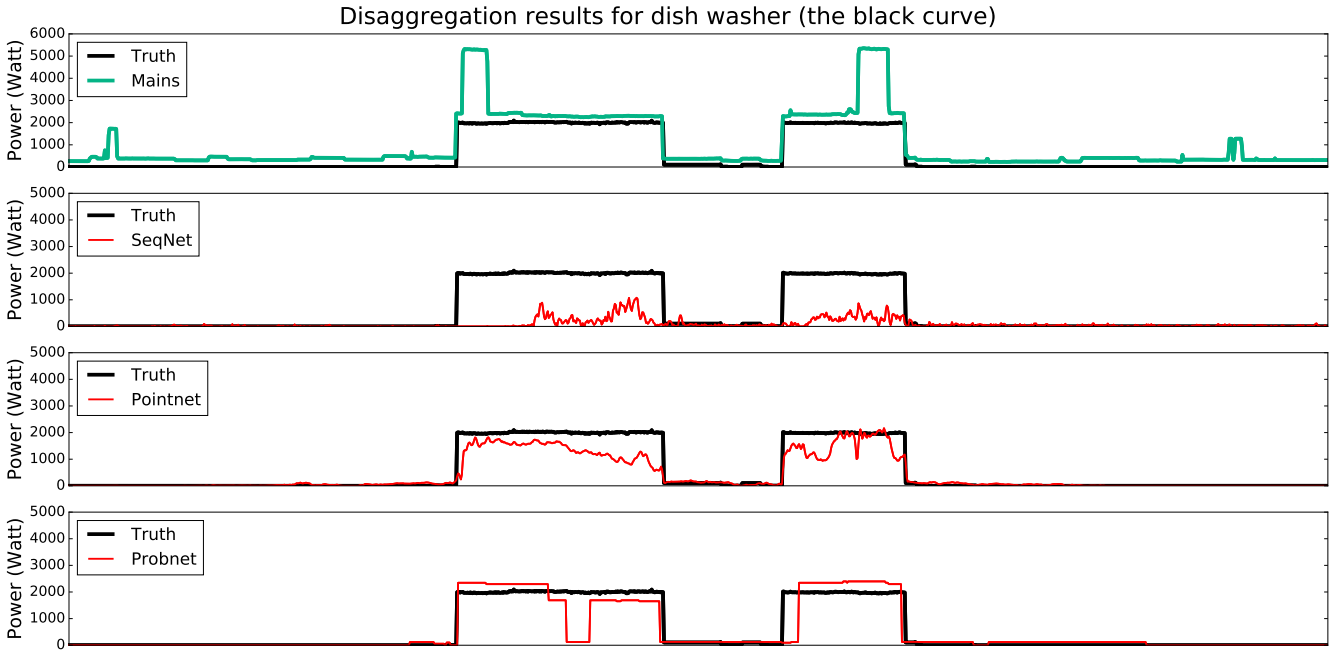

Figure 6: The example disaggregation results for dish washer.

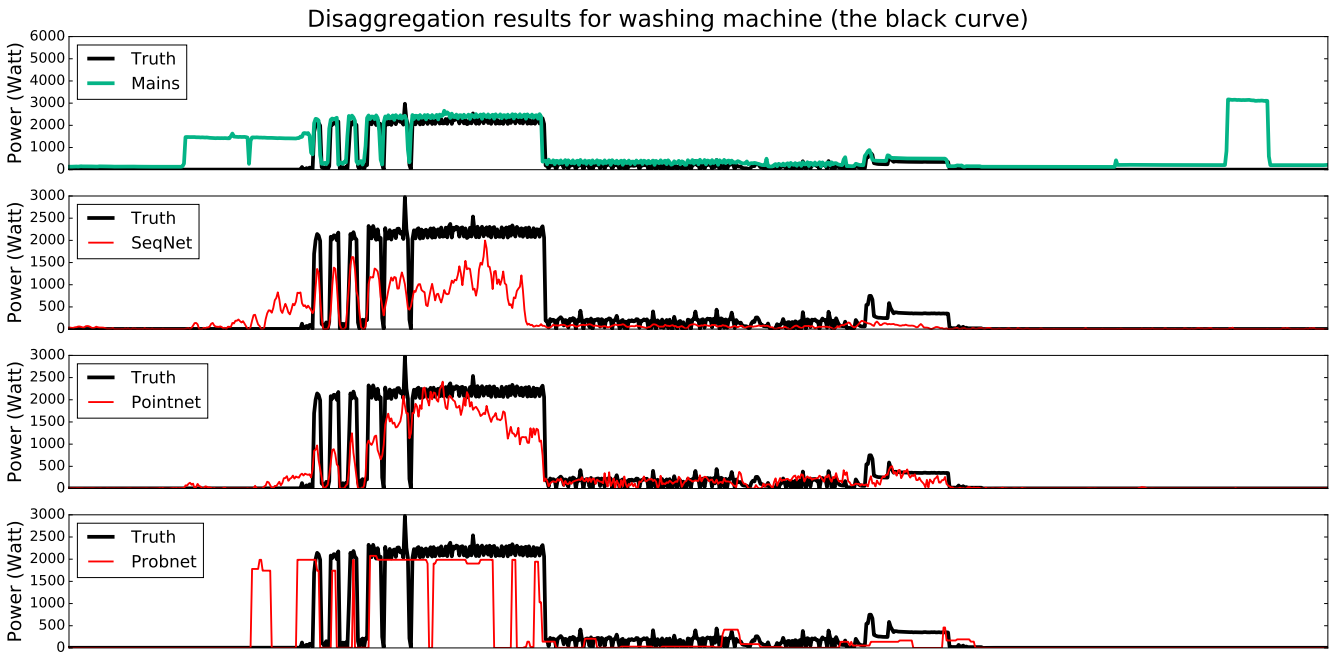

Figure 7: The example disaggregation results for washing machine.